arXiv:1702.06001v1 [physics.data-an] 1 Feb 2017

Optimal design of experiments by combining coarse and fine measurements

Alpha A. Lee* and Michael P. Brenner
School of Engineering and Applied Sciences and Kavli Institute of Bionano Science and Technology, Harvard University, Cambridge, MA 02138, USA

Lucy J. Colwell†
Department of Chemistry, University of Cambridge, CB2 1EW, Cambridge, UK

In many contexts it is extremely costly to perform enough high quality experimental measurements to accurately parameterize a predictive quantitative model. However, it is often much easier to carry out experiments that indicate whether a particular sample is above or below a given threshold. Can many such binary or "coarse" measurements be combined with a much smaller number of higher resolution or "fine" measurements to yield accurate models? Here, we demonstrate an intuitive strategy, inspired by statistical physics, wherein the coarse measurements identify the salient features of the data, while fine measurements determine the relative importance of these features. We illustrate our strategy by considering the problem of solubility prediction for small organic molecule from their 2D molecular structure.

A large class of scientific questions asks whether the values of dependent variables can be accurately predicted given sets of input variables and training data. Classical statistics shows that this is possible with sufficiently many high resolution measurements, though the cost of performing the required number of experimental measurements can be prohibitive. Nonetheless, in many settings, it can be straightforward to evaluate whether a sample measurement is above or below a certain threshold.

Examples of this problem abound in disparate fields. For example, predicting the solubility of organic molecules is a fundamental question in physical chemistry [1]. Although accurate measurement of the aqueous solubility of molecules is extremely challenging [2], measuring whether a molecule is soluble at a particular concentration is comparatively simple. Similarly, in the drug discovery workflow, biochemical assays that determine whether a molecule binds to a given receptor are relatively simple compared to accurate quantitative measurements of protein-ligand binding affinity [3]. In protein biophysics, a key challenge is to predict the effect of amino acid changes on protein phenotype. Here, threshold measurements are naturally provided by the set of homologous protein sequences that have similar phenotypic properties [4–8]. In contrast, experimentally measuring the biophysical properties is much more difficult. A related problem, crucial to HIV vaccine development, is to predict the fitness landscape of the virus given HIV sequences obtained from patients; again collecting patient samples is much easier than measuring the fitness directly [9, 10].

Despite the ubiquity of this problem, to our knowledge there is no principled method for combining numerous binary ("coarse") measurements with fewer quantitative ("fine") measurements to produce a predictive model. Although regression approaches can account for a prior estimate of sample error [11], this is not the same as combining two qualitatively distinct forms of data to build a more accurate model.

In this Letter, we introduce an intuitive method that combines coarse and fine measurements, inspired by the methods of statistical physics. We then apply this approach to the problem of predicting the solubility of chemical compounds. We interpret our method as deriving from data an accurate parameterization of an Ising model that relates the independent variables to the output variable.

To fix ideas, we assume each sample is characterized by a vector of p properties $\mathbf{f}_i \in \mathbb{R}^p$. We are given binary data for numerous samples of which N_+ (N_-) are above (below) some given threshold. In addition, we are given $\mathbf{y} \in \mathbb{R}^M$, quantitative measurements for M samples. These measurements could be biological fitness, binding affinity, solubility etc. We can define data matrices $R_\pm \in \mathbb{R}^{N_\pm \times p}$ for samples above/below the threshold, with columns normalized to have zero mean and unit variance. Intuitively, if there are combinations of the p properties that are always present in samples above or below the threshold, then these properties should be good predictors of the measurement. Such persistent correlations can be identified from the eigendecomposition of each sample covariance matrix $C_\pm$

$$\begin{aligned} C_\pm &= \frac{1}{N_\pm} R_\pm^T R_\pm \\ &= \sum_{i=1}^{N_\pm} \lambda_i^\pm \mathbf{u}_i^\pm \otimes \mathbf{u}_i^\pm, \end{aligned} \tag{1}$$

where $\{\lambda_i^\pm\}$, $\{\mathbf{u}_i^\pm\}$ are the eigenvalues and eigenvectors. Each $\mathbf{u}_i^\pm$ identifies a particular combination of the p properties that explains a fraction $\lambda_i^\pm / \sum_i \lambda_i^\pm$ of the data variance [11]. The matrix $C_\pm$ is an unbiased estimator of the corresponding true covariance matrix, which describes the relationships between variables in the distribution from which samples are drawn. This raises the

question of how the quality of this estimator depends on data sampling, and specifically whether the sample eigenvalues and eigenvectors are equally reliable. For example, one may preferentially assay samples with certain features (easier to obtain, or measure etc.), causing $C_\pm$ to have an eigenvector localized on these features with large corresponding eigenvalue, even though these features are uncorrelated with the output variable. For protein sequences, a natural source of such spurious correlation is phylogeny induced by evolutionary dynamics [12, 13].

Here we propose that overall the eigenvectors $\mathbf{u}_i^\pm$ computed from $C_\pm$, the covariance matrix of binary data measurements, in fact reliably identify data features, but the significance of these features as estimated by the corresponding eigenvalues $\lambda_i^\pm$ is severely corrupted by imperfect sampling. Later in this Letter we justify this ansatz with ideas from statistical physics, showing that this characterization is accurate for a large class of problems. Our ansatz implies a natural strategy is to use the quantitative measurements to determine the significance of each feature. In particular, we posit a general bilinear model

$$y_i = \mathbf{h}^T \mathbf{f}_i + \mathbf{f}_i^T J \mathbf{f}_i + \epsilon_i, \tag{2}$$

where $\mathbf{h}$ is the variable-specific effect, J captures the coupling between variables, and $\epsilon_i \sim N(0, \sigma)$ models random error. There are p parameters in $\mathbf{h}$ and $p(p-1)/2$ parameters in J. In principle those parameters could be estimated using linear least squares regression if one had $M \gg p(p+1)/2$ quantitative measurements. However, it is costly to perform a large number of detailed measurements, so we turn instead to the features identified from $C_\pm$. We pose the ansatz

$$J = \sum_{i=1}^{\hat{p}_+} c_i^+ \mathbf{u}_i^+ \otimes \mathbf{u}_i^+ + \sum_{i=1}^{\hat{p}_-} c_i^- \mathbf{u}_i^- \otimes \mathbf{u}_i^-, \tag{3}$$

where $\hat{p}_\pm \leq p$, since eigenvectors with low eigenvalues typically reflect noise due to finite sampling [14–16]. Our ansatz reflects the hypothesis that the eigendecomposition of $C_\pm$ captures variable-variable correlations. If the number of significant eigenvectors is much smaller than the number of variables, random matrix theory provides a rigorous way to determine $\hat{p}_\pm$ [14–19]; this case will be discussed in detail later. Relaxing this assumption, it is necessary to include all eigenvectors and determine the parameters $\mathbf{h}$, c^+ and c^- by regressing against the the few quantitative measurements available. We note that the ansatz (3), reduces the number of variables to $p + \hat{p}_+ + \hat{p}_-$.

To illustrate our strategy, we turn to the concrete example of solubility prediction. We work with solubility data for 1144 organic molecules reported in [20], the largest freely accessible dataset. It is well known that the aqueous solubility of a chemical compound relates to the presence of different chemical groups [21] — chemical groups that interact favorably (unfavorably) with water increase (or decrease) solubility. To describe the chemical groups in each molecule, we concatenate the Avalon Fingerprint [22], the MACCS Fingerprint [23], and the 1024-bit Morgan6 Fingerprint [24], using the package `rdKit` [25], to produce a vector of 1703 variables for each molecule. To generate binary data we label molecules with solubility $> 10^{-2}$ mol/L as soluble, and those with solubility $< 10^{-4}$ mol/L as insoluble. We then diagonalize the resulting matrices $C_\pm$ after eliminating variables that are zero for more that 1% of the dataset. For simplicity, we take $\hat{p}_\pm = \min\{p, N_\pm\}$, since it is not known *a priori* whether the number of significant eigenvectors is small compared to the number of variables. We randomly partition the dataset into a validation set (10%) and training set, and estimate $\{\mathbf{h}, \mathbf{c}^+, \mathbf{c}^-\}$ using linear regression with L_2 penalty on the coefficient to prevent overfitting (ridge regression), with the weight of the L_2 penalty estimated using 10 fold cross validation on the training set.

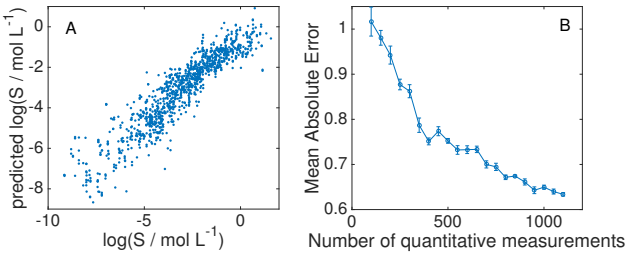

FIG. 1. Combining coarse and fine measurements accurately predicts solubility. (A) Out-of-sample solubility prediction has a mean absolute error of 0.62. The estimate is arrived at by analyzing 10 random partitions of the data into training and verification sets. (B) The mean absolute error as a function of the number of quantitative measurements given to the model. Error bars are obtained from averaging over 10 realizations.

Figure 1A shows that out-of-sample prediction for the validation set has a mean absolute error (MAE) of 0.62, comparable to state of the art models based on convolutional neural network [26, 27], though our model is significantly more parsimonious. Importantly, the prediction accuracy of our method remains robust when the number of quantitative measurements is reduced (Figure 1B). The error increases by less than a factor of 2 when the sample size decreases by an order of magnitude — confirming that our approach exploits the numerous binary measurements available to construct an accurate model using comparatively few quantitative measurements.

To understand why our heuristic strategy is successful, we consider a model problem where data is generated according to Eqn. (2), which is the maximum entropy model used in fields ranging from neuroscience [28, 29] to sociophysics [30] and is analogous to the Ising model.

We thus interpret the dependent variable as an "energy", noting that the logarithm of the solubility is indeed proportional to the solvation energy. The interaction matrix J can be decomposed into a sum of outer products of eigenvectors $\boldsymbol{\zeta}_i$ (Hopfield patterns [31]), and eigenvalues E_i (Hopfield energies) as

$$J = \sum_{i=1}^{m} E_i \boldsymbol{\zeta}_i \otimes \boldsymbol{\zeta}_i. \tag{4}$$

Furthermore, to model the binary features used in solubility prediction, we make the assumption that the independent variable is a vector of ± 1.

To simulate binary measurements we randomly draw samples from the uniform distribution, evaluate Eqn. (2) to determine the energy of each sample, and retain those samples that fall below a certain energy. Consider an interaction matrix J with $p = 100$, and $m = 3$ randomly generated patterns. To fix ideas, for the remainder of this Letter we let all patterns be attractive with $(E_1, E_2, E_3) = (-30, -25, -20)$, $\mathbf{h} = 0$ and $\epsilon_i = 0$. Using this model, we generate 5000 random vectors, take the $N = 500$ samples with lowest energy, and compute the eigendecomposition of the resulting correlation matrix.

Figure 2A shows that the eigenvalue distribution of the sample correlation matrix C follows the Marčenko-Pastur distribution expected for a random matrix [32],

$$\rho(\lambda) = \frac{\sqrt{\left[\left(1+\sqrt{\gamma}\right)^2 - \lambda\right]_+ \left[\lambda - \left(1-\sqrt{\gamma}\right)^2\right)]_+}}{2\pi\gamma\lambda} \tag{5}$$

where $\gamma = p/N$, except for three distinct eigenvalues. Figure 2B shows that those statistically significant eigenvectors are indeed the Hopfield patterns that we put in. Therefore, the large eigenvectors of the correlation matrix of the binary dataset correspond to eigenvectors of J. Note that the random matrix theory framework applies because $m \ll p$, i.e. the signal is low rank compared to the noise. If the signal were high rank, all eigenvectors must be included and their significance determined by regression against fine measurements, as in the solubility example discussed above.

Turning to eigenvalues, in this model which features *uniform* sampling, we find that the eigenvalues are proportional to the Hopfield energy E_i. This allows us to "clean" the correlation matrix, by taking the Marčenko-Pastur distribution as the null hypothesis, and using the q eigenvectors above the Marčenko-Pastur threshold to construct a rank q approximation J_{MP} of the correlation matrix (here $q = 3$). Figure 2C shows that J_{MP} accurately reconstructs J, and allows accurate prediction of the energy of particular states (Figure 2D). Analogously, the eigendecomposition of C_+ allows one to recover *repulsive* patterns with positive Hopfield energies (see Supplemental Information).

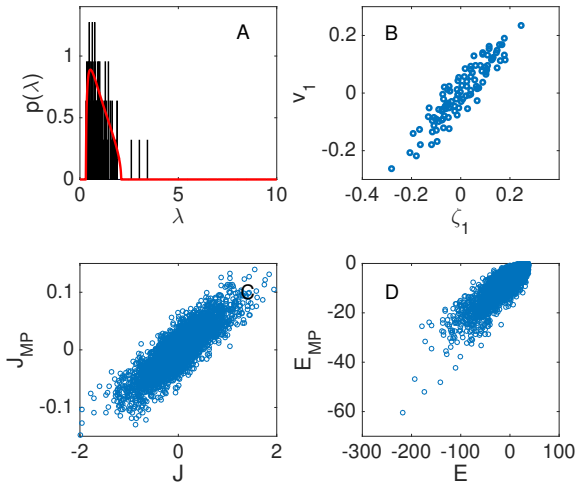

FIG. 2. Hopfield patterns can be recovered from threshold sampling. (A) Histogram of eigenvalues agrees with the Marčenko-Pastur distribution (red curve) save for three significant eigenvalues. (B) The top eigenvector is the lowest energy Hopfield pattern; the other eigenvectors are shown in Supplemental Information. Random matrix cleaning allows us to successfully (C) recover the coupling matrix J and (D) predict Hopfield energies.

Since the Hopfield patterns are energy minima, taken together Figure 2A-D imply that the probability of the system visiting a particular basin under uniform sampling is proportional to the energy of that minima. Therefore, the hypervolume of each energy basin is proportional to the basin depth. This fact can be derived by noting that Eqn. (2) is a quadratic form, so the Hessian matrix is a constant. Therefore, all local minima have the same mean curvature. Given that low lying energy minima are wide, we can extract the position of energy minima in the space of input variables by studying the correlation structure of the binary dataset. We note that the correlation between basin hypervolume and basin depth appears in many complex physical systems beyond the Ising model [33–35].

A lingering question is whether our inference procedure is robust to the energy threshold. To test this, we consider m Hopfield patterns, chosen as eigenvectors of a symmetrized $p \times p$ Gaussian random matrix, with the Hopfield energy chosen to be Gaussian distributed with mean 10 and unit variance. We draw 10000 samples randomly and compute the correlation matrix with the lowest energy N samples. Figure 3 shows that our method is robust: the correlation coefficient between J and J_{MP} is large and constant for a wide range of thresholds and number of Hopfield patterns. The question of how many energy minima can be recovered from the binary data, as well as an interpretation of our result based on the entropy of the energy basins, is discussed in the Supple-

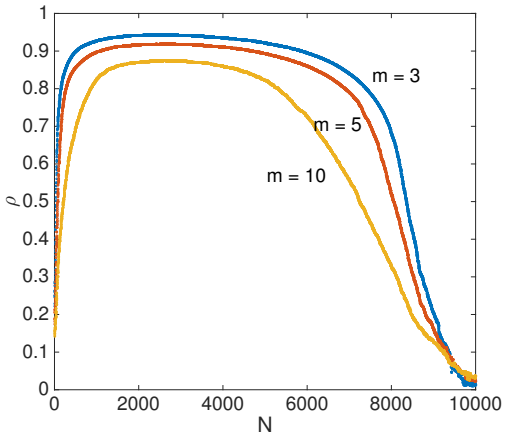

FIG. 3. Hopfield inference with random matrix cleaning is robust to the energy threshold. Here ρ is the Pearson correlation coefficient between the entries in J and $J_{\rm MP}$. The results are computed by averaging over 50 realizations.

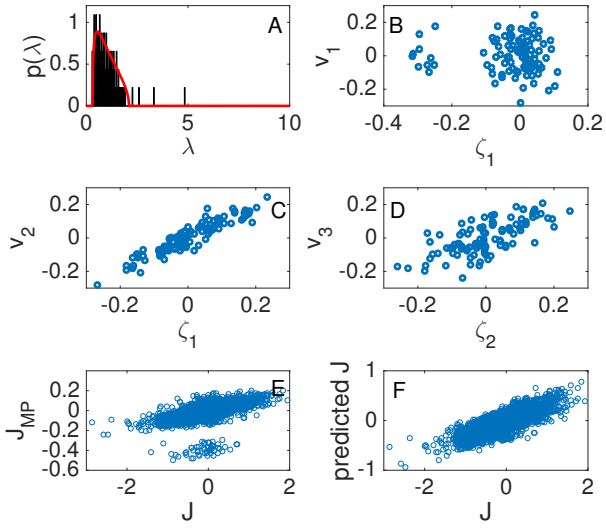

FIG. 4. Few quantitative measurements enable J to be inferred accurately for stratified datasets. (A) There are now four significant eigenvectors but still only three Hopfield patterns in the model. (B)-(D) The top eigenvector is uncorrelated with all Hopfield patterns, and the Hopfield patterns are demoted to the second to fourth significant eigenvectors. (E) Random matrix cleaning does not recover the coupling matrix. (F) The coupling matrix can be recovered by incorporating quantitative data and using Eqn. (3).

mental Information.

We now turn to consider two common scenarios that break the assumptions made so far, stratified sampling and more complex energy landscapes, to find out how they affect the eigenvectors and eigenvalues.

Stratified sampling: Thus far we assumed that the sampling is uniform before thresholding. However, in many settings, the sampling is biased, and some particular input variables are not well-sampled. To model this effect, we draw 5000 random samples, but freeze the first 5 variables to $+1$ for the first 2500 samples and the last 5 variables to -1 for the remaining 2500 samples. We then evaluate the energy, and take again the lowest $10^{\rm th}$ percentile. Figure 4A-D shows that the frozen variables introduce sample-sample correlations, and now there are 4 significant eigenvectors with the first Hopfield pattern demoted to the second largest eigenvector (Figure 4C). As such, the informative eigenvectors are still present in $J_{\rm MP}$, but the eigenvalues are misplaced.

Perhaps unsurprisingly, naïve random matrix cleaning does not recover J (Figure 4E), since there is no *a priori* reason to discard the first eigenvector unless we know beforehand the Hopfield patterns. We need additional information – for which we turn to the quantitative measurements — to accurately recover the Hopfield energies. Figure 4F shows that an additional 500 quantitative measurement allow us to recover the coupling matrix using ridge regression and the ansatz Eqn. (3).

Complex energy landscapes: The geometric property that the depth of an energy minima is related to its hypervolume is not universal to all energy landscapes [36, 37]. A natural question is whether the significant eigenvectors and eigenvalues in the correlation matrix of samples below/above an energy threshold allow us to infer features of a complex energy landscapes. We consider a landscape that comprises a sum of Gaussians

$$H(\mathbf{f}) = \sum_i E_i \exp\left(-E_i^2(\mathbf{f}\cdot\boldsymbol{\zeta}_i)^2\right). \tag{6}$$

This landscape has the property that the depth of each energy minima, E_i (located at $\boldsymbol{\zeta}_i$), is *inversely* proportional to its width $1/E_i$. As above, we let $(E_1, E_2, E_3) = (-30, -25, -20)$ and generate Hopfield patterns by diagonalising a symmetrized Gaussian random matrix. We draw 5000 samples and threshold to find the lowest 500 samples in energy. Figure 5 shows that there are again three significant eigenvectors above the Marčenko-Pastur threshold, but the lowest energy Hopfield pattern is demoted to the third eigenvector, while the highest energy Hopfield pattern is promoted to the top eigenvector. This is expected: the eigenvalue corresponding to each minimum is proportional to the number of samples near that minimum, i.e. the basin volume, which in this case is not proportional to basin depth. However, the eigenvectors still indicate the locations of the energy minima. Therefore a natural strategy is to use $\mathbf{f}\cdot\mathbf{u}_i$, the overlaps between the sample vector and each eigenvector, as inputs to a general nonlinear function such as an artificial neural network.

In conclusion, we develop a general strategy, grounded in statistical physics, which integrates coarse and fine

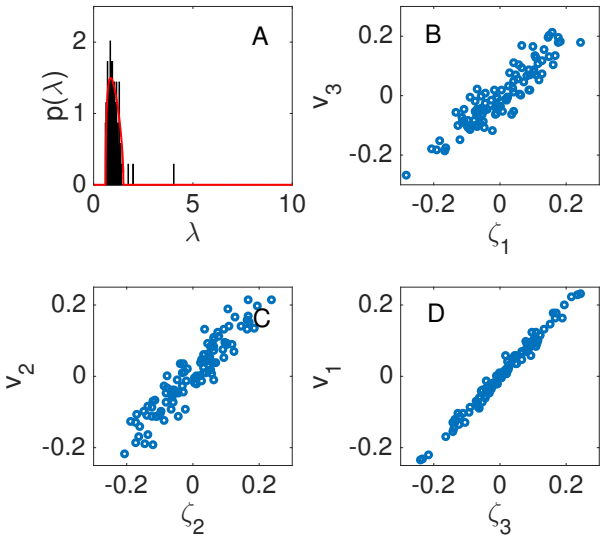

FIG. 5. For energy landscapes where basin depth is not proportional to basin width, the eigenvectors indicate the locations of energy minima but the eigenvalues are awry. (A) There are three significant eigenvectors above the Marčenko-Pastur threshold. (B) - (D) The top eigenvector is correlated with the highest energy minimum, and the last significant eigenvector is correlated with the lowest energy minimum.

measurements to yield a predictive model. Since coarse measurements are often significantly less costly to obtain, our strategy provides a new avenue for experiment design. Although our Letter only considered an Ising-type model, the fact that the eigenvectors of the correlation matrix of coarse measurements point toward energy minima suggests a natural way to integrate our result into more complex non-linear models.

The authors thank R Monasson for insightful discussions. AAL acknowledges the support of the George F. Carrier Fellowship. LJC acknowledges a Next Generation fellowship, and a Marie Curie CIG [Evo-Couplings, Grant 631609]. MPB is an investigator of the Simons Foundation, and acknowledges support from the National Science Foundation through DMS-1411694.

* alphalee@g.harvard.edu
† ljc37@cam.ac.uk

[1] A. Llinas, R. C. Glen, and J. M. Goodman, Journal of Chemical Information and Modeling **48**, 1289 (2008).
[2] D. S. Palmer and J. B. Mitchell, Molecular Pharmaceutics **11**, 2962 (2014).
[3] N. Malo, J. A. Hanley, S. Cerquozzi, J. Pelletier, and R. Nadon, Nature Biotechnology **24**, 167 (2006).
[4] F. Morcos, N. P. Schafer, R. R. Cheng, J. N. Onuchic, and P. G. Wolynes, Proceedings of the National Academy of Sciences **111**, 12408 (2014).
[5] M. Figliuzzi, H. Jacquier, A. Schug, O. Tenaillon, and M. Weigt, Molecular Biology and Evolution **33**, 268 (2015).
[6] T. A. Hopf, J. B. Ingraham, F. J. Poelwijk, M. Springer, C. Sander, and D. S. Marks, Nature Biotechnology (2017).
[7] P. Barrat-Charlaix, M. Figliuzzi, and M. Weigt, Scientific Reports **6** (2016).
[8] R. M. Levy, A. Haldane, and W. F. Flynn, Current Opinion in Structural Biology **43**, 55 (2017).
[9] K. Shekhar, C. F. Ruberman, A. L. Ferguson, J. P. Barton, M. Kardar, and A. K. Chakraborty, Physical Review E **88**, 062705 (2013).
[10] A. L. Ferguson, J. K. Mann, S. Omarjee, T. Ndungu, B. D. Walker, and A. K. Chakraborty, Immunity **38**, 606 (2013).
[11] C. M. Bishop, *Pattern recognition and Machine Learning*, Vol. 128 (Springer, 2007).
[12] J. Y. Dutheil, Briefings in Bioinformatics **13**, 228 (2012).
[13] B. Obermayer and E. Levine, New Journal of Physics **16**, 123017 (2014).
[14] L. Laloux, P. Cizeau, J.-P. Bouchaud, and M. Potters, Physical Review Letters **83**, 1467 (1999).
[15] V. Plerou, P. Gopikrishnan, B. Rosenow, L. A. N. Amaral, and H. E. Stanley, Physical Review Letters **83**, 1471 (1999).
[16] Z. Bai and J. W. Silverstein, *Spectral analysis of large dimensional random matrices*, Vol. 20 (Springer, 2010).
[17] J. Bun, R. Allez, J.-P. Bouchaud, and M. Potters, IEEE Transactions on Information Theory **62**, 7475 (2016).
[18] A. A. Lee, M. P. Brenner, and L. J. Colwell, Proceedings of the National Academy of Sciences **113**, 13564 (2016).
[19] J. Bun, J.-P. Bouchaud, and M. Potters, Physics Reports **666**, 1 (2017).
[20] J. S. Delaney, Journal of Chemical Information and Computer Sciences **44**, 1000 (2004).
[21] R. Skyner, J. McDonagh, C. Groom, T. van Mourik, and J. Mitchell, Physical Chemistry Chemical Physics **17**, 6174 (2015).
[22] P. Gedeck, B. Rohde, and C. Bartels, Journal of Chemical Information and Modeling **46**, 1924 (2006).
[23] J. L. Durant, B. A. Leland, D. R. Henry, and J. G. Nourse, Journal of Chemical Information and Computer Sciences **42**, 1273 (2002).
[24] D. Rogers and M. Hahn, Journal of Chemical Information and Modeling **50**, 742 (2010).
[25] "RDKit: Open-source cheminformatics," `http://www.rdkit.org`.
[26] D. K. Duvenaud, D. Maclaurin, J. Iparraguirre, R. Bombarell, T. Hirzel, A. Aspuru-Guzik, and R. P. Adams, in *Advances in Neural Information Processing Systems* (2015) pp. 2224–2232.
[27] S. Kearnes, K. McCloskey, M. Berndl, V. Pande, and P. Riley, Journal of Computer-aided Molecular Design **30**, 595 (2016).
[28] E. Schneidman, M. J. Berry, R. Segev, and W. Bialek, Nature **440**, 1007 (2006).
[29] S. Cocco, S. Leibler, and R. Monasson, Proceedings of the National Academy of Sciences **106**, 14058 (2009).
[30] E. D. Lee, C. P. Broedersz, and W. Bialek, Journal of Statistical Physics **160**, 275 (2015).
[31] J. J. Hopfield, Proceedings of the National Academy of Sciences **79**, 2554 (1982).
[32] V. A. Marčenko and L. A. Pastur, Mathematics of the

USSR-Sbornik **1**, 457 (1967).
[33] J. P. Doye, D. J. Wales, and M. A. Miller, The Journal of Chemical Physics **109**, 8143 (1998).
[34] J. P. Doye and C. P. Massen, The Journal of Chemical Physics **122**, 084105 (2005).
[35] C. J. Pickard and R. Needs, Journal of Physics: Condensed Matter **23**, 053201 (2011).
[36] D. J. Wales, M. A. Miller, and T. R. Walsh, Nature **394**, 758 (1998).
[37] D. Wales, *Energy landscapes: Applications to clusters, biomolecules and glasses* (Cambridge University Press, 2003).